

On the Mixing Time of Geographical Threshold Graphs

Andrew Beveridge*, Milan Bradonjić†

Abstract

We study the mixing time of random graphs in the d -dimensional toric unit cube $[0, 1]^d$ generated by the geographical threshold graph (GTG) model, a generalization of random geometric graphs (RGG). In a GTG, nodes are distributed in a Euclidean space, and edges are assigned according to a threshold function involving the distance between nodes as well as randomly chosen node weights, drawn from some distribution. The connectivity threshold for GTGs is comparable to that of RGGs, essentially corresponding to a connectivity radius of $r = (\log n/n)^{1/d}$. However, the degree distributions at this threshold are quite different: in an RGG the degrees are essentially uniform, while RGGs have heterogeneous degrees that depend upon the weight distribution. Herein, we study the mixing times of random walks on d -dimensional GTGs near the connectivity threshold for $d \geq 2$. If the weight distribution function decays with $\mathbb{P}[W \geq x] = O(1/x^{d+\nu})$ for an arbitrarily small constant $\nu > 0$ then the mixing time of GTG is $O(n^{2/d}(\log n)^{(d-2)/d})$. This matches the known mixing bounds for the d -dimensional RGG.

Keywords: Geographical threshold graphs, Mixing time, Random geometric graphs.

1 Introduction

In recent years, we have witnessed the development of numerous approaches to study the structure of large real-world technological and social networks, and to optimize processes on these networks. Large networks, such as the Internet, World Wide Web, phone call graphs, infections disease contacts and financial transactions, have provided new challenges for modeling and analysis [8]. As an example, Web graphs may have billions of nodes and edges, which implies that processing and extracting information on these large sets of data, is ‘hard’ [1]. Extensive theoretical and experimental research has been done in web-graph modeling, attempting to capture both the structure and dynamics of the web graph [9, 2, 23, 22, 26, 5, 3, 7, 17].

In general, a particularly fertile approach is to consider the network as an instance of an ensemble, arising from a suitable random generative model. Since the seminal papers on the evolution of uniform random graph model [20, 21], many other models have been proposed to better capture the structure seen in real-world networks, which are systematically covered in [19]. One straightforward example is the random geometric graph (RGG) model, where nodes are placed uniformly at random in a Euclidean space and edges are placed between any two nodes within a threshold distance. For further study of RGGs, see the monograph by Penrose [29]. The RGGs have the advantage

*Department of Mathematics, Statistics and Computer Science, Macalester College, Saint Paul, MN 55105, USA. E-mail: abeverid@macalester.edu.

†Mathematics of Networks and Communications, Bell Laboratories, Alcatel-Lucent, 600 Mountain Avenue, Murray Hill, NJ 07974, USA. E-mail: milan@research.bell-labs.com. This work was done while the author was at Los Alamos National Laboratory.

of describing many aspects of systems such as sensor networks, while avoiding unnecessary details. However, they fail to capture heterogeneity in the network.

Geographical threshold graphs (GTGs), introduced in [28], are a generalization of RGGs. Heterogeneity in the network is provided via a richer stochastic model that nevertheless preserves much of the simplicity of the RGG model. GTGs assign to nodes both a location and a weight. The weight may represent a quantity such as transmission power in a wireless network or influence in a social network. Edges are placed between two nodes if a symmetric function of their weights and the distance between them exceeds a certain threshold [13].

Structural properties of GTGs, such as connectivity, clustering coefficient, degree distribution, diameter, existence and absence of the giant component, chromatic number have been recently analyzed [11, 12, 14]. These properties are not merely of theoretical importance, but also play an important role in applications. In communication networks, connectivity implies the ability to reach all parts of the network. In packet routing, diameter gives the minimal number of hops needed for transmission between two arbitrary nodes. In the case of epidemics, the existence or absence of the giant component controls whether the epidemic spreads or is contained. When treating the node colors as the different radio channels or frequencies, the chromatic number gives the minimal number of channels needed so that neighboring radios do not interfere with each other.

Herein, we consider random walks on GTGs near the connectivity threshold. Random walks (or more formally, Markov chains) on large networks have many applications. For example, random walks model the spread of disease or the dispersion of information [10]. The *mixing time* of a random walk is the expected number of random steps that are required to guarantee that the current distribution is close to the stationary distribution. Mixing times are an essential tool in both theory and practice: for example, see the recent survey of Diaconis [18] on Markov chain Monte Carlo methods. In [6], the authors derived the mixing time of exponential random graphs, a model extensively used in sociology, showing that the mixing time of the *Glauber dynamics* was $\Theta(n^2 \log n)$ for the unimodal *Gibbs distribution*, and exponential for the multimodal case. For the definitions of the Glauber dynamics and Gibbs distribution see [6].

Upper bounds on the mixing time for RGG at the connectivity threshold have been well studied. For the 2-dimensional RGG, Avin and Ercal [4] showed that the mixing time is $O(n)$. More recently, Cooper and Frieze [16] proved that for $d \geq 3$, the mixing time of a d -dimensional RGG is $\tilde{O}(n^{2/d})$ (in this notation, the logarithmic factors are suppressed). In this paper, we study the mixing times of random walks on d -dimensional GTGs near the connectivity threshold, where $d \geq 2$. We prove that when the node weight distribution decays sufficiently quickly, the mixing time is $O(n^{2/d}(\log n)^{(d-2)/d}) = \tilde{O}(n^{2/d})$, which matches the mixing bounds for RGG. This result is formulated more precisely as Theorem 2 in the next section.

2 The GTG model and the mixing time

The GTG model is constructed from of a set of n nodes placed independently uniformly at random into the unit cube in \mathbf{R}^d . A non-negative weight w_i , taken randomly and independently from a continuous probability distribution function $f(w) : \mathbf{R}^+ \rightarrow \mathbf{R}^+$, is assigned to each node v_i for $i \in \{1, 2, \dots, n\}$. For nodes i, j at distance r_{ij} , the edge $\{i, j\}$ exists if and only if the following connectivity relation is satisfied:

$$G(w_i, w_j)D(r_{ij}) \geq \theta_n. \quad (1)$$

Here $G(w_i, w_j)$ is the interaction strength between nodes, $D(r)$ is a decreasing function of r and θ_n is a given threshold parameter that depends on the size n of the network. The interaction strength

$G(w_i, w_j)$ is usually taken to be symmetric and either multiplicatively or additively separable, i.e., in the form of $G(w_i, w_j) = g(w_i)g(w_j)$ or $G(w_i, w_j) = g(w_i) + g(w_j)$. We use $D(r) = r^{-s}$ where $s > 0$, which is a typical attenuation in the path-loss model in wireless communications [13].

Some basic results have already been shown, including the expected degree of a node with given weight w , when the nodes are distributed uniformly over a unit space [28, 13]. In both the multiplicative and additive cases of $G(w_i, w_j)$, questions of diameter, connectivity, and topology control have been addressed [13].

Here we restrict ourselves to nodes distributed uniformly over $[0, 1]^d$. For analytical simplicity we take the space to be the d -dimensional toric unit cube $[0, 1]^d$. Our connectivity relation uses an additive interaction strength $G(w_1, w_2) = w_1 + w_2$ and a decay function $D(r) = r^d$, so that nodes i, j are adjacent when

$$\frac{w_i + w_j}{r_{ij}^d} \geq \theta_n. \quad (2)$$

This connectivity relation identifies a d -dimensional sphere of influence for each vertex.

We assume that our weight distribution $f(w)$ has finite mean and finite variance. Let the cumulative density function (cdf), for the distribution of node weights $f(w)$, be

$$F(x) = \mathbb{P}[W \leq x] = \int_0^x f(w) dw. \quad (3)$$

The argument in [12] characterizing the degrees of a GTG for a 2-dimensional GTG is easily generalized to dimension d (only the leading constant changes). For any threshold $\theta_n = O(n)$ and any weight distribution such that $\mathbb{P}[W \geq x] = O(1/x^{1+\epsilon})$ for an arbitrarily small constant $\epsilon > 0$, the degree distribution of a node v with weight w follows the binomial distribution

$$\deg(v|w) \sim \text{Bin}(n-1, p(w)), \quad (4)$$

where $p(w) = \frac{\Upsilon_d}{\theta_n}(w + \mu)$, $\mu = \mathbb{E}[W]$ is the expected node weight and Υ_d is the volume of the unit ball in d dimensions:

$$\Upsilon_d = \frac{\pi^{d/2}}{\Gamma(d/2 + 1)} = \begin{cases} \pi^k/k!, & d = 2k \text{ even,} \\ 2^d k! \pi^k/d!, & d = 2k + 1 \text{ odd.} \end{cases}$$

Herein, we assume that $\mathbb{P}[W \geq x] = O(1/x^{d+\nu}) = o(1/x^2)$ which ensures that the weight distribution has finite mean and finite variance.

We now highlight the differences between the GTG and RGG models. The main characteristic of the GTG model is its tunable topology. By changing the input parameters $f(w)$ and θ_n , one can obtain graphs with different structural properties. For example, we can generate an RGG for a desired degree distribution, while the degrees of RGG are always uniform. The major distinction between our analysis for GTGs and the analysis for RGGs, lies in addressing the following two issues: (i) two spatially close nodes in a GTG are not necessarily connected, since they may both have very low weight; (ii) two distant nodes in a GTG are not necessarily disconnected, since one of them can carry heavy weight. These two issues never happen in the RGG model, and represent a challenge in our analysis.

Theorem 5.3 of [11] characterizes the connectivity threshold for a GTG in 2 dimensions. We list the changes to the proof to generalize to d -dimensional GTGs. We tile the unit space $[0, 1]^d$ into $\Theta(n/\log n)$ cubes of equal volume (as opposed to 2-dimensional squares). We use the connectivity relation equation (2) for the d -dimensional space and connectivity radius $r_{\text{conn}} = (\log(\alpha n)/(\alpha n \Upsilon_d))^{1/d}$

instead, with the constant $\alpha \in (0, 1)$ specified in Theorem 5.3 of [11]. With these modifications, we obtain the proof of the connectivity threshold for a d -dimensional GTG, given by the following theorem.

Theorem 1. *Let G be a GTG in the d -dimensional toric unit cube $[0, 1]^d$ with threshold function $\theta_n = cn / \log n$ where the constant $c < \sup_{\alpha \in (0, 1)} \alpha F^{-1}(1 - \alpha) / 4$. Then G is connected whp¹.*

In essence, the proof of Theorem 1 consists of two parts. First, one shows that the αn nodes of highest weight are connected. Next, one shows that the remaining $(1 - \alpha)n$ nodes are connected to first set. This partition into *high weight nodes* and *low weight nodes* will also be useful herein.

This paper gives an upper bound on the mixing time for a simple random walk on a geographical threshold graph $G = (V(G), E(G))$ at the connectivity threshold, provided that the weight distribution decays at an adequate rate. A simple random walk on a graph G (cf. [27]) consists of a sequence of vertices $(w_0, w_1, \dots, w_t, \dots)$ such that for $t \geq 0$, $\mathbb{P}[w_{t+1} = j \mid w_t = i]$ is $1 / \deg(i)$ if $\{i, j\} \in E(G)$ and 0 otherwise. Let π denote the *stationary distribution* of this random walk, so that $\pi_k = \deg(k) / 2|E|$ for every node k . Supposing that i is the initial node of our random walk, let $P^t(i, \cdot)$ denote the distribution of the states at time t . The *variational distance* at time t is

$$\Delta_i(t) = \frac{1}{2} \sum_{j \in V} |P^t(i, j) - \pi_j|.$$

When G is not bipartite, we have $\lim_{t \rightarrow \infty} \Delta_i(t) = 0$ for every $i \in V(G)$. The *mixing time* from node i measures how quickly $P^t(i, \cdot)$ converges to π . Explicitly, this mixing time from i is defined as

$$\tau_i(\delta) = \min\{t \mid \Delta_i(t') \leq \delta, \forall t' \geq t\}.$$

The mixing time of G is

$$\tau(\delta) = \max_{i \in V(G)} \tau_i(\delta).$$

We choose $\delta = 1/n$ as our desired distance from the stationary distribution. Our main result is given by the following theorem.

Theorem 2. *Let G be a connected GTG with threshold function $\theta_n = cn / \log n$ in the d -dimensional toric unit cube $[0, 1]^d$. If the weight distribution satisfies $\mathbb{P}[W \geq x] = O(1/x^{d+\nu})$ for $\nu > 0$, and*

$$c \leq \frac{1}{d+3} \min \left\{ \sup_{\alpha \in (0, 1/2]} \alpha F^{-1}(\alpha), \sup_{\alpha \in (1/2, 1)} (1 - \alpha) F^{-1}(\alpha) \right\},$$

then $\tau(1/n) = O(n^{2/d} (\log n)^{(d-2)/d})$, whp.

This mixing bound for GTG matches the best known mixing bound for RGG, and we believe that this equivalence is essentially correct. Intuitively, the αn high weight nodes of a GTG G contain a spanning subgraph G' that is an RGG. The mixing time of G corresponds to the mixing time of G' . There are some extremely long edges in $G \setminus G'$, but they seem to be too sparse to aid in mixing. At the same time, we find that the $(1 - \alpha)$ low weight nodes (with very short edges) do not slow down mixing. For technical reasons, we consider a weight decay of $\mathbb{P}[W \geq x] = O(1/x^{d+\nu})$. We conjecture that this equivalence continues to hold for $\mathbb{P}[W \geq x] = o(1/x^{1+1/d})$.

¹ We will use the notation “with high probability” and denote it as *whp*, meaning with probability $1 - o(1)$ as n tends to infinity.

The rest of the paper is organized as follows. In Section 3, we derive upper and lower bounds on the maximal weight. Consequently, we find upper and lower bounds on the node degrees, and show that the number of edges $|E(G)| = \Theta(n \log n)$. In Section 4 we construct a family of canonical paths for the GTG and then prove Theorem 2 in Section 5. In Section 6, we reflect on our results and explain why we believe that our result holds for more slowly decaying weight distributions. Finally, in A, we exemplify our results with two different weight functions. Our first example is the exponential weight distribution $f(w) = e^{-w}$, with cumulative density function $F(x) = 1 - e^{-x}$. Our second example is the Pareto distribution with cumulative density function $F(x) = 1 - x^{-\gamma}$, where $x \geq 1$ and $\gamma > d \geq 2$.

3 Node weights and node degrees in GTG

In this section, we determine the upper and lower bounds on the maximal weight W_{\max} in a geographical threshold graph $G = (V(G), E(G))$. Subsequently, we derive the upper and lower bounds on the degrees of the nodes in GTG near the connectivity threshold. Finally, we show that the number of edges $|E(G)| = \Theta(n \log n)$ for these connected GTGs.

We adopt the following notation for the remainder of the paper. We have a constant $\alpha \in (0, 1)$ and we fix small constants ϵ, ν so that

$$0 < \epsilon < \nu/2d. \quad (5)$$

Furthermore, we assume that there is a weight W_0 such that if $W \geq W_0$ then $\mathbb{P}[W \geq x] = O(1/x^{d+\nu})$. For brevity, we will state this as “ $\mathbb{P}[W \geq x] = O(1/x^{d+\nu})$.” Finally, we use $\kappa(n)$ to denote an arbitrarily slowly increasing function of n , that is $\kappa(n) = \omega(1)$.

The maximal weight W_{\max} satisfies $\mathbb{P}[W_{\max} \leq x] = F(x)^n$, since the weights are independently distributed. Consider a continuous weight distribution $f(w)$ with cdf $F(x)$. Our goal is to find two thresholds W_1, W_2 , such that $\mathbb{P}[W_{\max} \leq W_1] = o(1)$ and $\mathbb{P}[W_{\max} \geq W_2] = o(1)$. We can ‘invert’ $F(x)$ using the quantile function $F^{-1}(p) = \inf\{x \in \mathbf{R}^+ : p \leq F(x)\}$. Define

$$W_1 = F^{-1}\left(1 - \frac{\kappa(n)}{n}\right) \quad \text{and} \quad W_2 = F^{-1}\left(1 - \frac{1}{n \kappa(n)}\right).$$

We have $\mathbb{P}[W_{\max} \leq W_1] = (1 - \kappa(n)/n)^n \leq \exp(-\kappa(n)) \rightarrow 0$ and $\mathbb{P}[W_{\max} \geq W_2] = (1 - 1/n \kappa(n))^n \geq \exp(-1/\kappa(n) + 1/n \kappa(n)^2) \rightarrow 1$. In conclusion, the maximal weight W_{\max} satisfies

$$\lim_{n \rightarrow +\infty} \mathbb{P}[W_{\max} \in (W_1, W_2)] = 1. \quad (6)$$

See A for concrete examples of the calculation of the bounds on the maximal weight.

Let us determine the upper and lower bounds for the node degrees, keeping in mind that the weight distribution has finite mean and variance. We consider the GTG around the connectivity regime, as described in Theorem 1. The next result generalizes Lemma 3 in [16], which shows that all degrees of RGGs near the connectivity threshold are $\Theta(\log n)$ whp.

Lemma 3. *Let G be a connected GTG with threshold function $\theta_n = cn/\log n$. Whp, the nodes $v \in V(G)$ satisfy $\deg(v) \in I_{GTG}$ where*

$$I_{GTG} = \left[c_1 \log n, c_2 F^{-1}\left(1 - \frac{1}{n \kappa(n)}\right) \log n \right], \quad (7)$$

for any function $\kappa(n) = \omega(1)$ and for constants

$$c_1 = \frac{\mu\Upsilon_d}{c} \left(1 - \sqrt{\frac{2c}{\mu\Upsilon_d}}\right) \text{ and } c_2 = \frac{2\Upsilon_d}{c}.$$

The minimum degree of a GTG is $\Omega(\log n)$. The maximum degree depends upon the decay rate of the weight distribution: a slower decay rate results in larger maximum degree. In A, we calculate I_{GTG} for weight distributions with exponential decay and polynomial decay.

Proof. Using equation (4), we apply the Chernoff bound on the degree $\deg(v|w)$ of a node v with a given weight w :

$$\mathbb{P}\left[\deg(v|w) \leq (1 - \delta)\mathbb{E}[\deg(v|w)]\right] \leq \exp(-\mathbb{E}[\deg(v|w)]\delta^2/2).$$

Let $c_3 = 1 - \sqrt{\frac{2c}{\mu\Upsilon_d}} < 1$. By having $\mathbb{E}[\deg(v|w)] = (n-1)p(w) = \frac{\Upsilon_d(n-1)}{cn}(w + \mu)\log n$ and choosing $\delta = 1 - c_3/(1 + w/\mu)$, it follows

$$\mathbb{P}\left[\deg(v|w) \leq c_1 \log n\right] \leq \exp\left(-\frac{\mu\Upsilon_d}{c}(1 + w/\mu)\left(1 - \frac{c_3}{1 + w/\mu}\right)^2\left(1 - \frac{1}{n}\right)\log n\right). \quad (8)$$

Next, we specify conditions such that equation (8) is $o(1/n)$ for all $w \geq 0$ and sufficiently large n . For the sake of simplicity, let us denote $x = 1 + w/\mu \geq 1$, and consider the function

$$\phi(x) = \frac{\mu\Upsilon_d}{c}x\left(1 - \frac{c_3}{x}\right)^2.$$

The minimum of $\phi(x)$ is attained at $x = c_3$. Moreover, $\phi(x)$ is strictly decreasing on $(0, c_3)$ and strictly increasing on $(c_3, +\infty)$. Because $c_3 < 1$ and $\phi(1) = \frac{\Upsilon_d\mu}{c}(1 - c_3)^2 = 2$, we know that $\phi(x) \geq 2$ for $x \geq 1$. That is, equation (8) is $o(1/n)$, for $n \geq 3$. Thus, the degree distribution satisfies

$$\begin{aligned} \mathbb{P}\left[\deg(v) \leq c_1 \log n\right] &= \int f(w)\mathbb{P}\left[\deg(v|w) \leq c_1 \log n\right]dw \\ &= o\left(n^{-1} \int f(w)dw\right) = o(1/n). \end{aligned}$$

The union bound gives the lower bound on degree of the nodes in the graph.

We now obtain the upper bound. Equation (6) ensures $F^{-1}(1 - \kappa(n)/n) \leq W_{\max} \leq F^{-1}(1 - 1/n\kappa(n))$. Moreover, by the continuity of $F^{-1}(x)$, for any $\epsilon > 0$, there is sufficiently large $n = n(\epsilon)$, such that the upper and lower bounds on W_{\max} are arbitrarily close $0 \leq W_2 - W_1 \leq \epsilon$.

The degree of the node with maximal weight satisfies the binomial distribution $\text{Bin}(n-1, (\Upsilon_d/\theta_n)(W_{\max} + \mu))$, which is concentrated around its mean $(\Upsilon_d/c)(1 - 1/n)(W_{\max} + \mu)\log n$. Finally, the union bound gives the upper bound on the degrees. \square

We now partition the interval I_{GTG} of equation (7). We use this partition to calculate the number of edges $|E(G)|$ and again in Section 5 to bound the mixing time.

Define $h(x) = (1 - F(x))^{-1}$, or equivalently, $F(x) = 1 - 1/h(x)$. This is also equivalent to

$$h^{-1}(y) = F^{-1}(1 - 1/y).$$

By assumption we have $h(x) = \Omega(x^{d+\nu})$, so that $h^{-1}(x) = O(x^{1/(d+\nu)})$. Our first interval B contains the low weight nodes:

$$B = \{v \in V(G) \mid w_v \leq F^{-1}(1 - \alpha)\} = \{v \in V(G) \mid F(w_v) \leq 1 - \alpha\}. \quad (9)$$

Next, we partition the αn nodes with weights in $[F^{-1}(1 - \alpha), W_{\max}]$. By equation (6), $W_{\max} = F^{-1}(1 - 1/n\kappa(n)) = h^{-1}(n\kappa(n))$. Let

$$\begin{aligned} a_0 &= W_{\max} \\ a_k &= \max \left\{ F^{-1} \left(1 - a_{k-1}^{-(1+\epsilon)} \right), F^{-1}(1 - \alpha) \right\}, \text{ for } k \geq 1. \end{aligned}$$

The a_k are only defined until we reach $F^{-1}(1 - \alpha)$. Call this final index M . Our partition consists of the subintervals of the form $(a_k, a_{k-1}]$ for $1 \leq k \leq M$. Note that the indexing of our endpoints is the reverse of the standard convention.

Lemma 4. *The final index satisfies $M = o(\log n)$.*

Proof. We have $a_0 = W_{\max} = h^{-1}(n\kappa(n)) = O((n\kappa(n))^{1/(d+\nu)})$. Let $\beta = (1 + \epsilon)/(d + \nu) < 1$. By induction, for $0 < k < M$,

$$a_k = F^{-1} \left(1 - a_{k-1}^{-(1+\epsilon)} \right) = h^{-1} \left(a_{k-1}^{1+\epsilon} \right) = O \left(W_{\max}^{\beta^k} \right) = O \left((n\kappa(n))^{\beta^k/(d+\nu)} \right).$$

By definition, $a_M = F^{-1}(1 - \alpha)$. Let N be the smallest integer such that $(n\kappa(n))^{\beta^N/(d+\nu)} < F^{-1}(1 - \alpha)$. The right hand side is a constant, so $M = O(N) = O(\log \log(n\kappa(n))) = o(\log n)$. \square

For $1 \leq k \leq M$, let

$$A_k = \{v \in V(G) \mid w_v \in (a_k, a_{k-1}] \}. \quad (10)$$

The degree of $v \in A_k$ with weight w_v is

$$\deg(v) = O((1 + w_v) \log n) = O(a_{k-1} \log n). \quad (11)$$

We now show that the sizes of the A_k are all concentrated around their means.

Lemma 5. *Whp, for all $1 \leq k \leq M$, we simultaneously have $|A_k| = \Theta(n/a_{k-1}^{1+\epsilon})$.*

Proof. We have $\mathbb{P}[W > a_k] = 1/a_{k-1}^{1+\epsilon}$. Therefore $\mathbb{E}[|A_1|] = n/a_0^{1+\epsilon}$ and $\mathbb{E}[|A_k|] = n/a_{k-1}^{1+\epsilon} - n/a_{k-2}^{1+\epsilon} = \Theta(n/a_{k-1}^{1+\epsilon})$ for $2 \leq k \leq M$.

As for the concentration of these values, we consider the most delicate case of A_1 . We have $\mathbb{P}[v \in A_1] = \mathbb{P}[w_v > a_1] = a_0^{-(1+\epsilon)} = W_{\max}^{-(1+\epsilon)} := q_1$. We can consider A_1 as being generated from $\text{Bin}(n, q_1)$. The Chebyshev inequality gives

$$\mathbb{P}[|A_1| - nq_1| \geq nq_1] \leq \frac{1}{nq_1} = O \left(\frac{(n\kappa(n))^{(1+\epsilon)/(d+\nu)}}{n} \right) = O(n^{-1/2}).$$

Similarly, the probability that each remaining $|A_k|$ is more than twice its mean is $O(n^{-1/2})$. There are $M = o(\log n)$ such sets, so taking a union bound shows that all of them are concentrated whp. \square

See A for concrete examples of the partition of the nodes according to weight.

Lemma 6. $\sum_{k=0}^M a_k^{-\epsilon} = \Theta(1)$.

Proof. We use the general d'Alembert's convergence criterion: The sum of positive terms $\sum_{k=0}^{\infty} c_k$ is convergent if there exist a positive integer N and $\eta > 0$ such that $k > N$ guarantees $c_k/c_{k-1} < 1 - \eta$.

Let $b_0 = F^{-1}(1 - \alpha)$ and $b_k = h(b_{k-1})^{1/(1+\epsilon)}$ for $1 \leq k \leq M$. We have $b_k \leq a_{M-k-1}$ for $0 \leq k \leq M$. It follows that

$$\sum_{k=0}^M \frac{1}{a_k^\epsilon} \leq \frac{1}{a_M^\epsilon} + \sum_{k=0}^{\infty} \frac{1}{b_k^\epsilon} \leq 1 + \sum_{k=0}^{\infty} \frac{1}{b_k^\epsilon}.$$

and

$$\frac{1/b_k^\epsilon}{1/b_{k-1}^\epsilon} = \left(\frac{b_{k-1}}{b_k} \right)^\epsilon = \left(\frac{b_{k-1}}{h(b_{k-1})^{1/(1+\epsilon)}} \right)^\epsilon < 1 - \eta$$

for large k since $h(x) \gg x^{d+\nu} \gg x^{1+\epsilon}$. Therefore this sum converges to a constant (independent of n). \square

Lemma 7. *If G is a GTG at the connectivity threshold then the number of edges $|E(G)| = \Theta(n \log n)$ whp.*

Proof. Lemma 3 guarantees that all nodes have degree $\Omega(\log n)$ whp, so the number of edges $|E(G)| = \Omega(n \log n)$. As for the upper bound, whp

$$\begin{aligned} 2|E(G)| &= \sum_{v \in B} \deg(v) + \sum_{k=1}^M \sum_{v \in A_k} \deg(v) = O\left(\left(|B| + \sum_{k=1}^M |A_k|a_{k-1}\right) \log n\right) \\ &= O\left(\left(\alpha n + \sum_{k=1}^M \frac{n}{a_{k-1}^{1+\epsilon}} a_{k-1}\right) \log n\right) = O(n \log n), \end{aligned}$$

by Lemma 6. \square

4 Canonical paths for GTG

We employ canonical paths (as introduced in [25]) to calculate our bound on the mixing time. In this section, we construct the canonical paths for G , a connected GTG with threshold function $\theta_n = cn/\log n$. For every ordered pair of nodes $u, v \in V(G)$ we choose a canonical path γ_{uv} between them. We define

$$\begin{aligned} \rho &= \max_{e=\{x,y\} \in E(G)} \frac{1}{\pi(x)P(x,y)} \sum_{\gamma_{uv} \ni e} \pi(u)\pi(v)|\gamma_{uv}| \\ &= \max_{e=\{x,y\} \in E(G)} \frac{1}{2|E(G)|} \sum_{\gamma_{uv} \ni e} \deg(u)\deg(v)|\gamma_{uv}|, \end{aligned} \tag{12}$$

where $|\gamma_{uv}|$ is the length of the canonical path from u to v . As per [25] Proposition 12.1, the mixing time from node i satisfies

$$\tau_i(\delta) \leq \rho \left(\log \pi(i)^{-1} + \log \delta^{-1} \right), \tag{13}$$

where $\delta > 0$. We will set $\delta = 1/n$, so that $\tau_i(1/n) \leq \rho \left(\log \pi(i)^{-1} + \log n \right)$.

Let G be a GTG at the connectivity threshold $\theta_n = cn / \log n$ for $c < \sup_{\alpha \in (0,1)} \alpha F^{-1}(1-\alpha)/4$, as per Theorem 1. The proof of this result in [11] establishes the following two facts. First, a constant fraction of the nodes αn have weights greater than $F^{-1}(1-\alpha)$. We let $H(G) = \{i \in V(G) \mid w_i > F^{-1}(1-\alpha)\}$ denote this set of *high weight nodes* and let $L(G) = \{i \in V(G) \mid w_i \leq F^{-1}(1-\alpha)\}$ denote the complementary set of *low weight nodes*. Second, each high weight node is connected to every node within the *critical radius*

$$r_{\text{conn}} = \left(\frac{\log(\alpha n)}{\alpha n \Upsilon_d} \right)^{1/d}. \quad (14)$$

In other words, the induced subgraph on $H(G)$ contains a subgraph G' that is a connected RGG. Every vertex in $L(G)$ is adjacent to nodes in $H(G)$, so G is also connected.

We use the connected RGG subgraph G' to construct our canonical paths. Our construction for the canonical paths is similar to the one used in [16] to bound the mixing time of RGG. Compared to that result, our proof addresses a novel technical challenge: all of the degrees of a RGG are $\Theta(\log n)$, while the degrees of a GTG are heterogeneous. This leads to two challenges. First, when $u, v \in L(G)$, the intermediate nodes on the canonical path γ_{uv} must all be in $H(G)$. Therefore, we must choose these paths so that the intermediate nodes are evenly distributed among the high weight nodes. Second, many high weight nodes have degrees that are $\omega(\log n)$. We must ensure that their contribution to the sum in equation (12) does not lead to an increase in the mixing time, compared to a RGG.

We now describe the geometric scaffold for our canonical paths, as in [16]. Partition the unit cube into a toric grid of $[k]^d$ small cubes, where k is specified below. A set of canonical paths for the grid will act as the framework for our canonical paths for the GTG. Intuitively, we increment the entries (mod k) in succession. So first we increase a_1 until we achieve b_1 , then do the same for the second entry, and so on. The canonical path from (a_1, a_2, \dots, a_d) to (b_1, b_2, \dots, b_d) is

$$(a_1, a_2, \dots, a_d), (a_1 + 1, a_2, \dots, a_d), \dots, (b_1 - 1, a_2, \dots, a_d), (b_1, a_2, \dots, a_d), \\ (b_1, a_2 + 1, \dots, a_d), \dots, (b_1, b_2, \dots, b_d - 1), (b_1, b_2, \dots, b_d). \quad (15)$$

Each path has length at most dk . Note that we always increment the index by +1 (even if there is a shorter path).

While there are k^{2d} canonical paths, each edge appears in no more than k^{d+1} paths. Indeed, any path that includes the edge from $(i_1, \dots, i_t, \dots, i_d)$ to $(i_1, \dots, i_t + 1, \dots, i_d)$ must start at $(l_1, \dots, l_t, j_{t+1}, \dots, j_d)$ and end at $(j_1, \dots, j_t - 1, l'_t, \dots, l'_d)$ for some l_1, \dots, l_t and l'_t, \dots, l'_d . This results in t choices for l_1, \dots, l_t and $d - t + 1$ choices for l'_t, \dots, l'_d .

Before constructing the canonical paths for our GTG, we must prove two lemmas. Tile $[0, 1]^d$ into cubes S_i with side length $1/k = (c' \log n/n)^{1/d} = \Theta(r_{\text{conn}})$, where we state the conditions on constant c' later. We have a $[k]^d$ grid, whose cubes each have volume $c' \log n/n$. Let $H(S_i)$ and $L(S_i)$ denote the high weight and low weight nodes in S_i , respectively.

Lemma 8. *There exist constants $\beta_0, \beta_1 > 0$ such that whp every cube S satisfies $\beta_0 \log n \leq |L(S)| \leq \beta_1 \log n$ and $\beta_0 \log n \leq |H(S)| \leq \beta_1 \log n$.*

Proof. This proof is similar to the proof of Theorem 3 in [11], and will use the same notation. Let $B_i = |H(S_i)|$ and $R_i = |L(S_i)|$. In expectation, there are $\mathbb{E}[B_i] = \alpha c' \log n$ high weight nodes within

S_i . Using the lower and upper tail Chernoff bounds [15], it follows

$$\begin{aligned}\mathbb{P}[B_i \leq (1 - \delta_1)\mathbb{E}[B_i]] &\leq \exp\left(-\frac{\delta_1^2}{2}\mathbb{E}[B_i]\right), \\ \mathbb{P}[B_i \geq (1 + \delta_2)\mathbb{E}[B_i]] &\leq \exp\left(-\frac{\delta_2^2}{2(1 + \delta_2)}\mathbb{E}[B_i]\right).\end{aligned}$$

Fixing $\delta_2 \in [0, 1]$, we take $\delta_1 = \delta_2/\sqrt{1 + \delta_2}$, and thus $\delta_1 \in [0, 1]$. The number of high weighted nodes B_i within the cube S_i satisfies

$$\begin{aligned}\mathbb{P}[B_i \in (1 - \delta_1, 1 + \delta_2)\mathbb{E}[B_i]] &\geq 1 - 2\exp\left(-\frac{\delta_2^2}{2(1 + \delta_2)}\mathbb{E}[B_i]\right) \\ &= 1 - 2n^{-\alpha c' \delta_2^2 / (1 + \delta_2)}.\end{aligned}$$

By the union bound it follows:

$$\mathbb{P}\left[\bigcap_i \left\{B_i \in (1 - \delta_1, 1 + \delta_2)\mathbb{E}[B_i]\right\}\right] \geq \left(1 - 2n^{-\alpha c' \delta_2^2 / (1 + \delta_2)}\right)^{n/(c' \log n)}.$$

Taking the limit of the last expression as $n \rightarrow +\infty$, for $c' \geq (1 + \delta_2)/\alpha \delta_2^2$, we obtain the concentration on B_i , for each cube S_i :

$$\lim_{n \rightarrow \infty} \mathbb{P}\left[\bigcap_i \left\{B_i \in (1 - \delta_1, 1 + \delta_2)\alpha c' \log n\right\}\right] = 1.$$

The concentration on the number of low weight nodes R_i , within each cube S_i , follows analogously to the previous analysis. In expectation, there are $\mathbb{E}[R_i] = (1 - \alpha)c' \log n$ low weight nodes within S_i . Hence by the Chernoff tail bounds and the union bound we have

$$\lim_{n \rightarrow \infty} \mathbb{P}\left[\bigcap_i \left\{R_i \in (1 - \delta_1, 1 + \delta_2)(1 - \alpha)c' \log n\right\}\right] = 1.$$

Finally, we can guarantee the concentration of both the high and low weight nodes by taking $c' \geq \max\left\{\frac{1 + \delta_2}{\alpha \delta_2^2}, \frac{1 + \delta_2}{(1 - \alpha) \delta_2^2}\right\}$, that is,

$$c' \geq \begin{cases} \frac{1 + \delta_2}{\alpha \delta_2^2}, & \text{if } \alpha \in (0, 1/2) \\ \frac{1 + \delta_2}{(1 - \alpha) \delta_2^2}, & \text{if } \alpha \in [1/2, 1). \end{cases} \quad (16)$$

So the lemma holds with $\beta_0 = (1 - \delta_2/\sqrt{1 + \delta_2}) \min\{\alpha, 1 - \alpha\}$ and $\beta_1 = (1 + \delta_2) \max\{\alpha, 1 - \alpha\}$. \square

Two cubes are *adjacent* if they share a $(d - 1)$ -dimensional boundary.

Lemma 9. *Let S_i, S_j be adjacent cubes. The number of edges between $H(S_i)$ and $H(S_j)$ is $\Omega(\log^2 n)$ provided that $c \leq (d + 3)^{-d/2} \min\{\sup_{x \in (1/2, 1)}(1 - x)F^{-1}(x), \sup_{x \in (0, 1/2]}xF^{-1}(x)\}$.*

See A for explicit calculation of the constant c for two example weight distributions.

Proof. Consider any high weight node $u \in H(S_i)$ and any high weight node $v \in H(S_j)$, with the weights w_u and w_v , respectively. The distance r_{uv} between u and v is at most $\sqrt{d+3}(c' \log n/n)^{1/d}$. Indeed, the furthest points in two adjacent d -dimensional unit cubes are at distance $\sqrt{d+3}$ by the d -dimensional Pythagorean theorem.

Consider the connectivity relation

$$\frac{w_u + w_v}{r_{uv}^d} \geq \frac{F^{-1}(1-\alpha) + F^{-1}(1-\alpha)}{(d+3)^{d/2} c' \log n/n} = \frac{2F^{-1}(1-\alpha)}{(d+3)^{d/2} c'} \frac{n}{\log n}.$$

High weight nodes u, v are connected with probability one if $(w_u + w_v)/r_{uv}^d \geq \theta_n = cn/\log n$, which is guaranteed if $2F^{-1}(1-\alpha)/((d+3)^{d/2} c') \geq c$. Using equation (16), we require

$$c \leq \frac{2}{(d+3)^{d/2}} \frac{\delta_2^2}{1+\delta_2} \begin{cases} \alpha F^{-1}(1-\alpha), & \text{if } \alpha \in (0, 1/2) \\ (1-\alpha) F^{-1}(1-\alpha), & \text{if } \alpha \in [1/2, 1]. \end{cases}$$

Since $\delta_2 \in (0, 1)$ is arbitrary, with $\sup_{\delta_2 \in (0, 1)} \delta_2^2/(1+\delta_2) = 1/2$, the conditions combine to give

$$\begin{aligned} c &\leq (d+3)^{-d/2} \min \left\{ \sup_{\alpha \in (0, 1/2)} \alpha F^{-1}(1-\alpha), \sup_{\alpha \in [1/2, 1]} (1-\alpha) F^{-1}(1-\alpha) \right\} \\ &= (d+3)^{-d/2} \min \left\{ \sup_{x \in (1/2, 1)} (1-x) F^{-1}(x), \sup_{x \in (0, 1/2]} x F^{-1}(x) \right\}. \end{aligned}$$

□

We now employ a randomized procedure for choosing canonical paths. This procedure guarantees that no edge appears in more than $\Theta(k^{d+1}) = \Theta(r_{\text{conn}}^{-d-1}) = \Theta((n/\log n)^{(d+1)/d})$ canonical paths. Let us denote cubes as $S(a_1, \dots, a_d)$ for $1 \leq a_i \leq k$, where (a_1, \dots, a_d) is the location in the d -dimensional grid. By Lemma 8, we have both $\beta_0 \log n \leq |L(S(a_1, \dots, a_d))| \leq \beta_1 \log n$ and $\beta_0 \log n \leq |H(S(a_1, \dots, a_d))| \leq \beta_1 \log n$ for some constants $\beta_0, \beta_1 > 0$. For each cube $S = S(a_1, \dots, a_d)$, evenly partition $L(S)$ into sets $L_i(S)$, for $1 \leq i \leq \beta_0 \log n$. Each set has size at most $\lceil \beta_1/\beta_0 \rceil$ whp. Assign each set of low weight nodes to a distinct high weight node $h_i \in H(S)$, for $1 \leq i \leq \beta_0 \log n$. The high weight node h_i is the *high weight representative* for the nodes in $L_i(S)$. The key outcome of this assignment is that each $h \in H(S)$ represents a constant number of low weight nodes whp.

Consider any ordered pair of nodes (x, y) . Let $x \in S(a_1, \dots, a_d)$ and $y \in S(b_1, \dots, b_d)$. We choose a canonical path from x to y as follows. We use the toric grid to identify the sequence of cubes in the canonical path. Taking equation (15) as our framework, we consider cubes $S(a_1, a_2, \dots, a_d)$, $S(a_1+1, a_2, \dots, a_d), \dots, S(b_1, a_2, \dots, a_d)$, $S(b_1, a_2+1, \dots, a_d), \dots, S(b_1, b_2, \dots, b_d)$. For brevity, call these cubes S_0, S_1, \dots, S_t . If x is a low weight node, then set x_0 to be the high weight representative for x . If x is a high weight node, set $x_0 = x$. For $1 \leq i \leq t$, choose x_i to be a random high weight node in S_i . Set $x_t = y$ if y is a high weight node, otherwise using the high weight representative for y . Our canonical (x, y) path is $x, x_0, x_1, \dots, x_t, y$. If $x_0 = x$ or $x_t = y$, we remove the repeated node from the path. We have

$$\max_{u, v \in V(G)} |\gamma_{uv}| \leq d k + 2 = O\left((n/\log n)^{1/d}\right) \quad (17)$$

and furthermore, we can bound how often each edge appears in a canonical path.

Lemma 10. *Every edge in G appears in at most $O((n/\log n)^{(d+1)/d})$ canonical paths whp.*

Proof. Let Z_{xy} denote the number of times the edge xy is chosen. If both x and y are low weight nodes then $Z_{xy} = 0$ (we always move to a high weight node from a low weight node). When x is low weight and y is high weight, then the edge xy can only be used when x and y are in the same cube. In this case, the edge xy can only be chosen when the canonical path has x as one of its endpoints. Therefore $Z_{xy} \leq 2(n-1)$.

Now suppose that both x and y are high weight nodes. The edge xy is only used if x, y are in the same cube or in adjacent cubes. First, we consider high weight x, y in the same cube S . The edge xy will be used only by paths between $L_i(S) \cup \{x\}$ and $L_j(S) \cup \{y\}$. If one or both is not a high weight representative, then this edge will be used even fewer times. Therefore $Z_{xy} \leq (1 + \lceil \beta_1/\beta_0 \rceil)^2 = O(1)$ whp.

Consider $x \in S$ and $y \in T$ in adjacent cubes S, T . Let $\mathcal{P}_{xy} = \{\gamma_{uv} \mid xy \in \gamma_{uv}\}$. We consider four cases, according to the locations of u and v . If $u \in S$ and $v \in T$ then just as in the previous case, whp there are at most $(1 + \lceil \beta_1/\beta_0 \rceil)^2 = O(1)$ such paths.

The total number of paths γ_{uv} with $u \in S$ and $v \notin T$ is $O(n)$. Indeed, x must be the high weight representative of u , which gives $1 + \lceil \beta_1/\beta_0 \rceil = O(1)$ choices for u , and there are $O(n)$ end nodes for paths that start in S_1 . Similarly, the total number of paths with $u \notin S$ and $v \in T$ is $O(n)$.

The remaining case is when $u \notin S$ and $v \notin T$. Let Z'_{xy} denote the number of times the edge xy is chosen as the random edge from S to T for some canonical path. Let

$$\mathcal{P}(S, T) = \{\gamma_{uv} \mid \gamma_{uv} \text{ traverses from } S \text{ to } T \text{ and } u \notin S \text{ and } v \notin T\}.$$

Then whp

$$\mathbb{E}[Z'_{xy}] = \sum_{\gamma_{uv} \in \mathcal{P}(S, T)} \mathbb{P}(xy \in \gamma_{uv}) \leq r_{\text{conn}}^{-d-1} (\beta_1 \log n)^2 \frac{1}{(\beta_0 \log n)^2} = \left(\frac{\beta_1}{\beta_0}\right)^2 r_{\text{conn}}^{-d-1}.$$

Indeed, the number of canonical paths for the toric grid which pass from S to T is $O(r_{\text{conn}}^{-d-1})$. Each edge between cubes corresponds to at most $(\beta_1 \log n)^2$ toric paths in the GTG. Since both x and y are internal nodes of these paths, they were each chosen uniformly and independently with probability at most $1/(\beta_0 \log n)$.

Using the Chernoff bound for this binomial distribution,

$$\begin{aligned} \mathbb{P}[|Z'_{xy} - \mathbb{E}[Z'_{xy}]| \geq \eta \mathbb{E}[Z'_{xy}]] &\leq 2 \exp\left(-\frac{\eta^2}{3} \mathbb{E}[Z'_{xy}]\right) \\ &\leq 2 \exp\left(-\frac{\eta^2}{3} \left(\frac{\beta_1}{\beta_0}\right)^2 r_{\text{conn}}^{-d-1}\right) \\ &= 2 \exp\left(-\frac{\eta^2}{3} \left(\frac{\beta_1}{\beta_0}\right)^2 \left(\frac{n}{\alpha \Upsilon_d \log n}\right)^{(d+1)/d}\right) = o(1/e^n). \end{aligned}$$

The union bound now gives

$$\begin{aligned} &\mathbb{P}\left[\bigwedge_{x,y} (|Z'_{xy} - \mathbb{E}[Z'_{xy}]| \geq \eta \mathbb{E}[Z'_{xy}])\right] \\ &\leq \sum_{x,y} \mathbb{P}[|Z'_{xy} - \mathbb{E}[Z'_{xy}]| \geq \eta \mathbb{E}[Z'_{xy}]] \leq n^2 o(1/e^n) \rightarrow 0. \end{aligned}$$

Therefore whp, every edge between high weight nodes in adjacent cubes is used by $(1 \pm \eta)r_{\text{conn}}^{-d-1} = \Theta((n/\log n)^{(d+1)/d})$ canonical paths. In this case, $\mathbb{E}[Z_{xy}] = \Theta((n/\log n)^{(d+1)/d})$. \square

5 The mixing time for GTG

In this section, we prove that the mixing time for a d -dimensional GTG near the threshold for connectivity is $O(n^{2/d}(\log n)^{(d-2)/d})$ when $\mathbb{P}[W \geq x] = O(1/x^{d+\nu})$ for $\nu > 0$.

We have $|E(G)| = \Theta(n \log n)$ and $\max_{u,v \in V(G)} |\gamma_{uv}| = O((n/\log n)^{1/d})$ by Lemma 7 and equation (17), respectively. Substituting these values into equation (12) yields

$$\rho \leq O\left(\frac{1}{n^{(d-1)/d}(\log n)^{(d+1)/d}}\right) \left(\max_{e=\{x,y\} \in E(G)} \sum_{\gamma_{uv} \ni e} \deg(u) \deg(v)\right). \quad (18)$$

Fix an edge $e = \{x, y\}$ between high weight nodes in adjacent cubes, and define $\sigma(e) = \sum_{\gamma_{uv} \ni e} \deg(u) \deg(v)$. Suppose that these high weight cubes differ in the k th coordinate. Specifically, the cubes containing x and y are indexed by $(j_1, \dots, j_{k-1}, j_k, j_{k+1}, \dots, j_d)$ and $(j_1, \dots, j_{k-1}, j_k + 1, j_{k+1}, \dots, j_d)$, respectively. If $e \in \gamma_{uv}$ then u must be in an initial cube indexed by $(i_1, \dots, i_{k-1}, i_k, j_{k+1}, \dots, j_d)$, while v must be in a target cube indexed by $(j_1, \dots, j_{k-1}, \ell_k, \ell_{k+1}, \dots, \ell_d)$. Let $U_1 \subset [0, 1]^d$ correspond to the set of initial cubes, and let $U_2 \subset [0, 1]^d$ correspond to the set of target cubes. The volumes of these subsets are $\text{vol}(U_1) = (n/\log n)^{k/d}/(n/\log n) = (\log n/n)^{(d-k)/d}$ and $\text{vol}(U_2) = (n/\log n)^{(d-k+1)/d}/(n/\log n) = (\log n/n)^{(k-1)/d}$. Let V_1 (respectively V_2) be the set of nodes in U_1 (respectively U_2).

In order to bound $\sigma(e)$ we use the partition of the vertices into the sets B and A_1, A_2, \dots, A_M as defined in equations (9) and (10). We consider canonical paths containing e that run between every possible pair of sets in this partition. The following three technical lemmas require the weight decay of $\mathbb{P}[W \geq x] = O(1/x^{d+\nu})$.

Lemma 11. *For any $T \in \{B, A_1, A_2, \dots, A_M\}$, let $Z_i = |T \cap V_i|$ for $i = 1, 2$. Then $\mathbb{E}[Z_i] = |T| \text{vol}(U_i)$ and there exists a constant $C > 0$ (depending on T and i) such that*

$$\mathbb{P}[|Z_i - \mathbb{E}[Z_i]| \geq \mathbb{E}[Z_i]] = O\left(\exp(-Cn^{\nu/2d(d+\nu)})\right).$$

In other words, the probability that $Z_i \neq \Theta(\mathbb{E}[Z_i])$ decays at a super-polynomial rate.

Proof. We prove this concentration for $Z = A_M \cap V_1$, which is the most delicate case. The other cases follow similarly. First, by the independence of weights and location, we have $|A_M \cap V_1| \sim \text{Bin}(|A_M|, \text{vol}(U_1))$. Furthermore, by Lemma 5, $|A_M| = \Theta(n/a_{M-1}^{1+\epsilon})$. Therefore $A_M \cap V_1 \sim \text{Bin}(C'n/a_{M-1}^{1+\epsilon}, (\log n/n)^{(d-k)/d})$ for some $C' > 0$. The expected value is

$$\begin{aligned} \mathbb{E}[Z] &= \frac{C'n}{a_{M-1}^{1+\epsilon}} \left(\frac{\log n}{n}\right)^{(d-k)/d} \geq \frac{C'n}{W_{\max}^{1+\epsilon}} \left(\frac{\log n}{n}\right)^{(d-k)/d} \\ &\geq \frac{C'n}{(n\kappa(n))^{(1+\epsilon)/(d+\nu)}} \frac{(\log n)^{1/d}}{n^{(d-1)/d}} \\ &\geq C'n^{1/d-(1+\epsilon)/(d+\nu)} = \Omega\left(n^{\nu/2d(d+\nu)}\right) \end{aligned} \quad (19)$$

since $\epsilon < \nu/2d$ by equation (5) and $\kappa(n)$ grows arbitrarily slowly. Using the binomial Chernoff bound (cf [24], Corollary 2.3), we have

$$\mathbb{P}[|Z - \mathbb{E}[Z]| \geq \mathbb{E}[Z]] \leq 2 \exp\left(-\frac{1}{3}\mathbb{E}[Z]\right) = O\left(\exp(-Cn^{\nu/2d(d+\nu)})\right)$$

for some constant $C > 0$. □

For $T_1, T_2 \in \{B, A_1, A_2, \dots, A_M\}$, let $\lambda(T_1, T_2)$ denote the number of canonical paths from $T_1 \cap V_1$ to $T_2 \cap V_2$ that use edge e .

Lemma 12. *For $T_1, T_2 \in \{B, A_1, A_2, \dots, A_M\}$, we have*

$$\mathbb{E}[\lambda(T_1, T_2)] = \Theta\left(\frac{\mathbb{E}[|T_1 \cap V_1|] \cdot \mathbb{E}[|T_2 \cap V_2|]}{\log^2 n}\right)$$

and there exists a constant K (depending on T_1, T_2) such that

$$\mathbb{P}[\lambda(T_1, T_2) - \mathbb{E}[\lambda(T_1, T_2)] \geq \mathbb{E}[\lambda(T_1, T_2)]] = O\left(\exp(-Kn^{\nu/2d(d+\nu)})\right).$$

Proof. The proof is similar to the previous one. For $i = 1, 2$, let $Z_i = T_i \cap V_i$, so that $|Z_i| = \Theta(\mathbb{E}[Z_i])$ by Lemma 11. The distribution for $\lambda(T_1, T_2)$ is $\text{Bin}(|Z_1||Z_2|, K'/\log^2 n)$ for some $K' > 0$, and $|Z_i| = \Theta(\mathbb{E}[Z_i])$ for $i = 1, 2$ by Lemma 11.

Furthermore, $|Z_i| = \Omega(n^{\nu/2d(d+\nu)})$ by equation (19). Therefore $\mathbb{E}[\lambda(T_1, T_2)] = \Omega(n^{\nu/d(d+\nu)} / \log^2 n) = \omega(n^{\nu/2d(d+\nu)})$. The binomial Chernoff bound gives

$$\begin{aligned} \mathbb{P}[|\lambda(T_1, T_2) - \mathbb{E}[\lambda(T_1, T_2)]| \geq \mathbb{E}[\lambda(T_1, T_2)]] &\leq 2 \exp\left(-\frac{1}{3}\mathbb{E}[\lambda(T_1, T_2)]\right) \\ &= O\left(\exp(-Kn^{\nu/2d(d+\nu)})\right) \end{aligned}$$

for some constant $K > 0$. \square

The previous two lemmas show that these quantities are tightly concentrated around their means. A union bound shows that whp these quantities are concentrated for all candidate edges simultaneously. Indeed, there are $O(d \cdot (n/\log n)^{1/d})$ choices for adjacent cubes used in canonical paths, with $\Theta(\log^2 n)$ edges running between each pair. Recall that $M = o(\log n)$ by Lemma 4. For a fixed choice of adjacent cubes, there are $2(M+1) = o(\log n)$ choices for T in Lemma 11. Our union bound for these events involves $O((n/\log n)^{1/d} \log n)$ terms, each decaying at a super-polynomial rate. Therefore, all these events are concentrated whp. As for Lemma 12, there are $(M+1)^2 = o(\log^2 n)$ choices for (T_1, T_2) for each pair of adjacent cubes. This union bound is taken over all relevant edges between high adjacent cubes. The number of addends in this union bound is $O((n/\log n)^{1/d} (\log^2 n) (\log^2 n))$. Again, the concentrations from Lemma 12 are super-polynomial, so the union bound shows that all these quantities are concentrated simultaneously whp.

Lemma 13. *Our choice of canonical paths gives $\rho = O((\log n)^{2/d})$.*

Proof. We bound $\sigma(e) = \sum_{\gamma_{uv} \ni e} \deg(u) \deg(v)$. First consider the contributions from canonical paths from $B \cap V_1$ to $B \cap V_2$. Let $\sigma(B, B)$ denote the contribution of paths between low weight nodes to $\sigma(e)$.

Recall that if $v \in B$ then $\deg(v) = \Theta(\log n)$. By Lemmas 12, and 11, we have whp

$$\begin{aligned} \mathbb{E}[\sigma(B, B)] &= O\left(\frac{\lambda(B, B)}{\log^2 n} \cdot (\log^2 n)\right) = O(|B \cap V_1| \cdot |B \cap V_2|) \\ &= O\left(n \left(\frac{\log n}{n}\right)^{(d-k)/d} \cdot n \left(\frac{\log n}{n}\right)^{(k-1)/d}\right) \\ &= O\left(n^{(d+1)/d} (\log n)^{(d-1)/d}\right). \end{aligned}$$

We consider the contribution of paths between high weight nodes by using the partition A_1, A_2, \dots, A_M specified in equation (10). Let $\sigma(A_r, A_s)$ denote the contribution to $\sigma(e)$ for paths from A_r to A_s . Recall that, as per equation (11), if $v \in A_k$ then $\deg(v) = O(a_{k-1} \log n)$. Arguing similarly to the calculation above, the contribution to $\sigma(e)$ from paths between nodes in A_r and A_s , where $0 \leq r, s \leq M$, is whp

$$\begin{aligned}\sigma(A_r, A_s) &= O\left(\frac{\lambda(A_r, A_s)}{\log^2 n} \cdot a_r a_s \log^2 n\right) \\ &= O\left(\frac{n}{a_r^{1+\epsilon}} \left(\frac{\log n}{n}\right)^{(d-k)/d} \cdot \frac{n}{a_s^{1+\epsilon}} \left(\frac{\log n}{n}\right)^{(k-1)/d} a_r a_s\right) \\ &= O\left(\frac{n^{(d+1)/d} (\log n)^{(d-1)/d}}{a_r^\epsilon a_s^\epsilon}\right).\end{aligned}$$

Next we consider the paths between low weight and high weight nodes: whp

$$\begin{aligned}\sigma(B, A_s) &= O\left(\frac{\lambda(B, A_s)}{\log^2 n} \cdot a_s \log^2 n\right) \\ &= O\left(n \left(\frac{\log n}{n}\right)^{(d-k)/d} \cdot \frac{n}{a_s^{1+\epsilon}} \left(\frac{\log n}{n}\right)^{(k-1)/d} a_s\right) \\ &= O\left(\frac{n^{(d+1)/d} (\log n)^{(d-1)/d}}{a_s^\epsilon}\right).\end{aligned}$$

and similarly, $\sigma(A_r, B) = O(n^{(d+1)/d} (\log n)^{(d-1)/d} / a_r^\epsilon)$. Putting these estimates together, whp every edge e between cubes satisfies

$$\begin{aligned}\sigma(e) &= \sigma(B, B) + \sum_{j=1}^M (\sigma(B, A_j) + \sigma(A_j, B)) + \sum_{j=1}^M \sum_{k=1}^M \sigma(A_j, A_k) \\ &= O\left(n^{(d+1)/d} (\log n)^{(d-1)/d} \left(1 + 2 \sum_{i=1}^M \frac{1}{a_i^\epsilon} + \sum_{j=1}^M \sum_{k=1}^M \frac{1}{a_j^\epsilon a_k^\epsilon}\right)\right) \\ &= O\left(n^{(d+1)/d} (\log n)^{(d-1)/d} \left(1 + 2 \sum_{i=1}^M \frac{1}{a_i^\epsilon} + \left(\sum_{j=1}^M \frac{1}{a_j^\epsilon}\right)^2\right)\right) \\ &= O\left(n^{(d+1)/d} (\log n)^{(d-1)/d}\right)\end{aligned}$$

where the last equality follows from Lemma 6.

Finally, by equation (18) we have whp $\rho = O((n/\log n)^{2/d})$. □

Proof of Theorem 2. Equation (13) gives $\tau_x(\delta) \leq \rho (\log \pi(x)^{-1} + \log \delta^{-1})$. The previous lemma ensures that $\rho = O((n/\log n)^{2/d})$. Meanwhile, we have $\pi(x) = \deg(x)/2|E(G)| = \Omega(\log n/(n \log n)) = \Omega(1/n)$ by Lemma 3 and Lemma 7. In summary, $\tau_x(\delta) = O(n^{2/d} (\log n)^{(d-2)/d})$ for $\delta = 1/n$. □

6 Conclusion

We have shown that if the weight distribution of a d -dimensional GTG satisfies $\mathbb{P}[W > x] = O(1/x^{d+\nu})$, then its mixing time is $O(n^2/d(\log n)^{(d-2)/d})$. This matches the known bounds for RGG. Our proof uses a spanning subgraph among the αn high weight nodes to create a scaffold for canonical paths. In constructing these paths, our proof treats all high weight nodes identically, ignoring the particularly large reach of the highest weight nodes. We did try to take advantage of these hub nodes, but found that they were too sparse to leverage for canonical paths.

One might wonder whether this is a shortcoming of the method of canonical paths, rather than a reflection on the characteristics of GTG. However, initial investigations using conductance to bound mixing (as in [4]) suggests the same conclusion. Of course, using conductance introduces its own set of technical challenges due to the heterogeneous nature of the degrees.

For technical reasons, we assumed that the weight distribution decayed as $\mathbb{P}[W \geq x] = 1/x^{d+\nu}$. We believe that the equivalence of mixing for GTG and RGG extends to weight distributions with slower decay. In particular, we conjecture that these mixing time of a GTG with weight decay $\mathbb{P}[W \geq x] = O(1/x^\gamma)$ matches that of RGG when $\gamma \geq 1 + 1/d$, and that GTG mixes faster when $1 < \gamma < 1 + 1/d$.

Acknowledgments

The first author was supported in part by NSA Young Investigator Grant H98230-08-1-0064. The second author was supported in part through the Laboratory Directed Research and Development Program, and Center for Nonlinear Studies at Los Alamos National Laboratory.

References

- [1] ABELLO, J., PARDALOS, P. M., AND RESENDE, M. G. C., Eds. *Handbook of massive data sets*. Kluwer Academic Publishers, Norwell, MA, USA, 2002.
- [2] AIELLO, W., BONATO, A., COOPER, C., JANSSEN, J. C. M., AND PRALAT, P. A spatial web graph model with local influence regions. In *WAW* (2007), pp. 96–107.
- [3] AIELLO, W., CHUNG, F., AND LU, L. A random graph model for massive graphs. In *STOC '00: Proceedings of the thirty-second annual ACM symposium on Theory of computing* (New York, NY, USA, 2000), ACM Press, pp. 171–180.
- [4] AVIN, C., AND ERCAL, G. On the cover time and mixing time of random geometric graphs. *Theor. Comput. Sci.* 380, 1-2 (2007), 2–22.
- [5] BARABÁSI, A.-L., AND ALBERT, R. Emergence of scaling in random networks. *Science* 286, 5439 (1999), 509 – 512.
- [6] BHAMIDI, S., BRESLER, G., AND SLY, A. Mixing time of exponential random graphs. In *49th Annual IEEE Symposium on Foundations of Computer Science* (2008), pp. 803–812.
- [7] BOLLOBÁS, B., RIORDAN, O., SPENCER, J., AND TUSNÁDY, G. The degree sequence of a scale-free random graph process. *Random Struct. Algorithms* 18, 3 (2001), 279–290.

[8] BONATO, A. A survey of models of the web graph. In *Combinatorial and Algorithm Aspects of Networking* (Heidelberg, 2005), vol. 3405 of *Lecture Notes in Computer Sciences*, Springer Berlin, pp. 159–172.

[9] BONATO, A., HADI, N., HORN, P., PRALAT, P., AND WANG, C. A dynamic model for on-line social networks. In *WAW* (2009), pp. 127–142.

[10] BOYD, S., GHOSH, A., PRABHAKAR, B., AND SHAH, D. Randomized gossip algorithms. *IEEE Transactions on Information Theory, Special issue of IEEE Transactions on Information Theory and IEEE ACM Transactions on Networking* 52, 6 (June 2006), 2508–2530.

[11] BRADONJIĆ, M., HAGBERG, A., AND PERCUS, A. G. The structure of geographical threshold graphs. *Internet Mathematics* 5, 1–2 (November 2008), 113–139.

[12] BRADONJIĆ, M., HAGBERG, A. A., AND PERCUS, A. G. Giant component and connectivity in geographical threshold graphs. In *Proceedings of the 5th Workshop on Algorithms and Models for the Web-Graph (WAW2007)* (2007), A. Bonato and F. R. K. Chung, Eds., vol. 4863 of *Lecture Notes in Computer Science*, Springer, pp. 209–216.

[13] BRADONJIĆ, M., AND KONG, J. Wireless ad hoc networks with tunable topology. In *Proceedings of the 45th Annual Allerton Conference on Communication, Control and Computing* (2007).

[14] BRADONJIĆ, M., MÜLLER, T., AND PERCUS, A. G. Coloring geographical threshold graphs. *Discrete Mathematics & Theoretical Computer Science* 12, 3 (2010), 103–114.

[15] CHERNOFF, H. A note on an inequality involving the normal distribution. *Ann. Probab.* 9, 3 (1981), 533–535.

[16] COOPER, C., AND FRIEZE, A. The cover time of random geometric graphs. *Proceedings of SODA* (2009), 48–57.

[17] COOPER, C., AND FRIEZE, A. M. A general model of undirected Web graphs. In *ESA '01: Proceedings of the 9th Annual European Symposium on Algorithms* (London, UK, 2001), Springer-Verlag, pp. 500–511.

[18] DIACONIS, P. The markov chain monte carlo revolution. *Bull. Amer. Math. Soc.* 46 (2009), 127–205.

[19] DURRETT, R. *Random Graph Dynamics*, 1st ed. Cambridge University Press, 2006.

[20] ERDŐS, P., AND RÉNYI, A. On random graphs I. *Publ. Math. Debrecen* 6 (1959), 290–297.

[21] ERDŐS, P., AND RÉNYI, A. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci.* 5 (1960), 17–61.

[22] FLAXMAN, A. D., FRIEZE, A. M., AND VERA, J. A geometric preferential attachment model of networks. *Internet Mathematics* 3, 2 (2007).

[23] FLAXMAN, A. D., FRIEZE, A. M., AND VERA, J. A geometric preferential attachment model of networks ii. In *WAW* (2007), pp. 41–55.

[24] JANSON, S., ŁUCZAK, T., AND RUCIŃSKI, A. *Random Graphs*. John Wiley & Sons, 2000.

- [25] JERRUM, M., AND SINCLAIR, A. The markov chain monte carlo method: an approach to approximate counting and integration. In *Approximation Algorithms for NP-hard Problems* (1996), D. S. Hochbaum, Ed., PWS Publishing.
- [26] KUMAR, R., RAGHAVAN, P., RAJAGOPALAN, S., SIVAKUMAR, D., TOMKINS, A., AND UPFAL, E. Stochastic models for the web graph. In *FOCS '00: Proceedings of the 41st Annual Symposium on Foundations of Computer Science* (Washington, DC, USA, 2000), IEEE Computer Society, p. 57.
- [27] LOVÁSZ, L. Random walks on graphs: a survey. In *Miklós, D., Sós, V.T., Szőnyi, T. (eds.) Combinatorics, Paul Erdős is Eighty. Vol II* (1996), pp. 355–397.
- [28] MASUDA, N., MIWA, H., AND KONNO, N. Geographical threshold graphs with small-world and scale-free properties. *Physical Review E* 71, 3 (2005), 036108.
- [29] PENROSE, M. D. *Random Geometric Graphs*. Oxford University Press, 2003.

A Characteristics of GTG for example weight distributions

We describe the relevant characteristics of GTGs for two different weight distributions: exponential decay and polynomial decay.

A.1 Exponential Weight Distribution

Our first example is the exponential weight distribution $f(w) = e^{-w}$ with cumulative density function $F(x) = 1 - e^{-x}$. Inverting the cdf gives $F^{-1}(x) = -\log(1 - x)$.

We first discuss the weights and degrees of the nodes in the GTG, as described in Section 3. As per equation (6), the maximum weight satisfies

$$W_{\max} \in [\log n - \kappa(n), \log n + \kappa(n)]$$

whp. Lemma 3 guarantees that whp all the node degrees are in the interval

$$I_{GTG} = [\Theta(\log n), \Theta(\log^2 n)] .$$

Next, we partition the interval I_{GTG} . By equation (9), the cutoff for low weight nodes is $F^{-1}(1 - \alpha) = \log \alpha^{-1}$. We partition the high weight nodes into disjoint subsets as specified in equation (10). Note that $F^{-1}(1 - x^{-(1+\epsilon)}) = \log(x^{1+\epsilon}) = \Theta(\log x)$. Therefore, ignoring leading constants, the sequence of endpoints (in descending order) is

$$(\log n, \log \log n, \log \log \log n, \dots, \log^{(M)} n)$$

where $M \leq \log^* n$, where the iterative logarithm function $\log^* n$ is the number of iterations of the \log function required to obtain a result less than 1.

Finally, we calculate upper bound on the constant c required by Lemma 9 in Section 4. For the exponential distribution, taking $c \leq ((d + 3)^{d/2} e)^{-1}$ is sufficient.

A.2 Pareto Weight Distribution

We now give a parallel characterization of our second example: a Pareto distribution with cumulative density function $F(x) = 1 - x^{-\gamma}$ where $\gamma > d \geq 2$. Inverting this cdf gives $F^{-1}(x) = (1-x)^{-1/\gamma}$.

We consider our results concerning node weights and node degrees. We have

$$W_{\max} \in \left[\left(\frac{n}{\kappa(n)} \right)^{1/\gamma}, (n\kappa(n))^{1/\gamma} \right]$$

whp by equation (6). By Lemma 3, whp all the node degrees are in

$$I_{GTG} = [\Theta(\log n), \Theta((n\kappa(n))^{1/\gamma} \log n)].$$

We separate the low weight nodes from the high weight nodes using the weight cutoff $F^{-1}(1-\alpha) = \alpha^{-1/\gamma}$. Next, we partition the high weight nodes as per equation (10). Note that $F^{-1}(1 - x^{-(1+\epsilon)}) = x^{(1+\epsilon)/\gamma}$. It follows that our sequence of endpoints is

$$((n\kappa(n))^{1/\gamma}, (n\kappa(n))^{\beta/\gamma}, (n\kappa(n))^{\beta^2/\gamma}, \dots, (n\kappa(n))^{\beta^M/\gamma})$$

where $\beta = (1+\epsilon)/\gamma < 1$. Here M is the smallest integer such that $(n\kappa(n))^{\beta^M/\gamma} \leq F^{-1}(1-\alpha) = \alpha^{-1/\gamma}$. The latter requirement is equivalent to the condition $M \geq (\log \log(n\kappa(n)) - \log \log(\alpha^{-1}))/\log(\beta^{-1})$.

Finally, we calculate an upper bound on the constant c required by Lemma 9 in Section 4. Since $F^{-1}(x) = (1-x)^{-1/\gamma}$, by Lemma 9 it follows $c \leq ((d+3)^{d/2} 2^{(\gamma-1)/\gamma})^{-1}$.